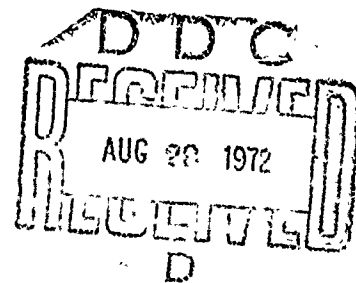
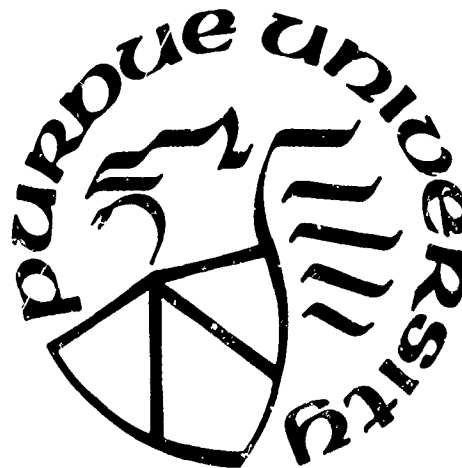


**PURDUE UNIVERSITY
SCHOOL OF ELECTRICAL ENGINEERING**

AD 747506



Lafayette, Indiana 47907

Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
U S Department of Commerce
Springfield VA 22151

Approved for public release;
distribution unlimited.

**BAYES DECISION RULES
BASED ON OBJECTIVE PRIORS[†]**

**PART I: FORMULATION AND APPLICATION
PART II: JUSTIFICATION**

T. L. Oberlin[‡] and R. L. Kashyap

**TR-EE 72-10
February 1972**

[†] This work was supported in part by the Air Force Office of Scientific Research under grant AFOSR-69-1776

[‡] Dr. Oberlin is presently with General Research Corporation, 1501 Wilson Blvd., Suite 700, Arlington, Virginia 22209.

DOCUMENT CONTAINS DATA - I & D

1. ORIGINATING ACTIVITY (Agency name)		2. PROJECT SECURITY CLASSIFICATION	
School of Electrical Engineering Purdue University Lafayette, Indiana 47907		UNCLASSIFIED	
3. REPORT TITLE			
BAYES DECISION RULES BASED ON OBJECTIVE PRIORS PART I: FORMULATION AND APPLICATION; PART II: JUSTIFICATION			
4. DESCRIPTIVE NOTES (Type of report and inclusive dates)			
Scientific Interim			
5. AUTHOR(S) (First name, middle initial, last name)			
T. L. Oberlin and R. L. Kashyap			
6. REPORT DATE	7. TOTAL NO. OF PAGES	7b. NO. OF REFS	
February 1972	27	10	
8a. CONTRACT OR GRANT NO.	9a. ORIGINATOR'S REPORT NUMBER(S)		
AFOSR-69-177b			
b. PROJECT NO.			
9769			
c. 61102F	9b. OTHER REPORT NO'S (Any other numbers that may be assigned this report)		
d. 681304	AFOSR - TR - 72 - 1593		
10. DISTRIBUTION STATEMENT			
Approved for public release; distribution unlimited.			
11. SUPPLEMENTARY NOTES		12. SPONSORING MILITARY ACTIVITY	
TECH OTHER		Air Force Office of Scientific Research(OM) 1400 Wilson Boulevard Arlington, Virginia 22209	
13. ABSTRACT			
<p>The problem of statistical decision making under uncertainty is considered. A Bayes approach based upon prior probabilities which are found using an objective inference technique developed by R. L. Kashyap is proposed as the basic solution procedure. The problem is formulated in a statistical decision theory format and the general solution technique outlined. Several examples are solved to illustrate its application.</p> <p>Using this inference technique, it is possible to have different priors for different experiments. A general decision criterion is formulated to handle these situations.</p>			

CONTENTS

<u>Section</u>		<u>Page</u>
	Part I: Formulation and Application	
I	INTRODUCTION	1
II	STATEMENT OF DECISION MAKING PROBLEM	3
III	GENERAL SOLUTION PROCEDURE	7
IV	EXAMPLE 1 - URN PROBLEM	12
V	EXAMPLE 2 - BALL PROBLEM	16
VI	SELECTION OF EXPERIMENT	21
VII	CONCLUSIONS	31
	REFERENCES	32
	Part II: Justification	
I	INTRODUCTION	1
II	SINGLE EXPERIMENT	1
III	MULTIPLE EXPERIMENTS	16
IV	CONCLUSIONS	25
	REFERENCES	27

BAYES DECISION RULES BASED ON OBJECTIVE PRIORS

PART I: FORMULATION AND APPLICATION

By T. L. Oberlin and R. L. Kashyap

Abstract

The problem of statistical decision making under uncertainty is considered. A Bayes approach based upon prior probabilities which are found using an objective inference technique developed by R. L. Kashyap is proposed as the basic solution procedure. The problem is formulated in a statistical decision theory format and the general solution technique outlined. Several examples are solved to illustrate its application.

Using this inference technique, it is possible to have different priors for different experiments. A general decision criterion is formulated to handle these situations.

I. INTRODUCTION

A good deal of current research is directed at systems having varying degrees of uncertainty. A common method of handling these uncertainties is to assume they can be characterized by random variables which are then assigned probability distributions.

Basically there are two schools of thought concerning this assignment procedure--the subjective and the objective. Using the subjective approach as outlined by Savage¹, the decision maker attempts to form a consistent set of preferences from which he obtains his prior distribution; note that two decision makers with the same information may have different priors. On the other hand, the objective approach attempts to present necessary or logical means of obtaining the priors. Using this approach, two decision makers having the same information will have the same priors, hence the term objective. Proponents of objective approaches range from Bernoulli (principle of insufficient reason) to Jeffreys², Jaynes³, Tribus⁴, and Kashyap^{5,6,7,8}.

In this paper a new approach to the problem of statistical decision making under uncertainty is presented. It uses Bayes decision rules based on prior probabilities found by applying the objective inference technique developed by Kashyap. The problem is formulated in a statistical decision theory format and the general solution technique outlined. Several examples are solved to illustrate its application.

The major emphasis in this part of the paper is on the formulation and use of the decision-making technique. Questions of justification are postponed until Part II.

The major result of this paper is the formulation of an objective Bayes approach to the problem of decision making under uncertainty. This study also points out the difficulties associated with using any inference technique whose inference varies with the experiment. A general approach which permits different prior probabilities is formulated to handle these situations.

II. STATEMENT OF THE DECISION MAKING PROBLEM

A. BASIC ELEMENTS

Assuming that the uncertainty in the decision problem can be represented by certain random variables--so called states of nature--the basic elements of the decision problem under uncertainty are as follows:

- θ : unknown state of nature, $\theta \in \Theta$, set of possible states of nature.
- a : actions available to decision maker, $a \in A$, set of potential actions.
- $L(\theta, a)$: loss function, $L(\theta, a)$ is a real-valued function defined $\Theta \times A$.
- e : experiment, $e \in E$, set of potential experiments.
- y : outcome of the experiment, $y \in Y_e$, set of possible outcomes of an experiment e .
- $d(y)$: decision rule, $d(y)$ selects an action from A for every possible outcome y .

The observations y are dependent on θ through the conditional probabilities $P(y|\theta)$, which are assumed to be known for all $y \in Y_e$, $\theta \in \Theta$. $P(y|\theta)$ is assumed to have a fixed set of values throughout the experiment; if it changes, so does the experiment.

For certain experiments it will be convenient to refer to a sequence of independent observations which will be denoted by $y(1), y(2), \dots, y(m+1)$.

Unless otherwise noted these variables will have the same distribution as y .

There also may be additional, testable prior information about θ other than its domain Θ . Information about θ is testable if, given any probability assignment for θ , it is (potentially) possible to verify whether or not this probability agrees with the information³. For example, it may be known that the average value of θ is greater than a certain given number or that the probability of a particular outcome is less than a given value. For the latter case, since $P(y|\theta)$ is known, this implies a restriction on the allowable probability distributions of θ .

Other types of loss functions considered are those which are functions of an outcome of the experiment and the action selected prior to the observance of this outcome. Their form is simply $L(y,a)$. However, in order to emphasize the difference in loss functions--which is essential to our approach, as we shall soon see--we shall use z instead of y to denote the observation when loss functions $L(z,a)$ are considered.

Both of these types of loss functions-- $L(\theta,a)$, $L(z,a)$ --shall be considered in this paper. In the former θ is the variable of interest and in the latter z is the variable of interest. In order to account for different experiments, observations, and decision rules, more notationally convenient forms for these loss functions will be used. If d_e represents a decision rule for experiment e , then these loss functions can be represented by $L(\theta,d_e)$ or $L(z,d_e)$ (see Ferguson⁹).

Primary consideration in this paper will be on the selection of single decision rules, d_e , and not on selection of sequences of decision rules.

B. DEFINITIONS OF BAYES DECISION RULE, RESTRICTED BAYES DECISION RULE

The decision making problem is essentially the selection of decision rules, d_e^* , which for a given experiment e minimize the expected value of the loss. Then the experiment which results in the minimum expected loss is selected as the optimal experiment e^* . (More will be said on this in Sec. VI.)

For a given probability distribution $P_e(\theta)$, the risk $r(P_e, d_e)$ is given by

$$r(P_e, d_e) = E_{\theta} R(\theta, d_e) = E_{\theta} [E_{y|\theta} L(\theta, d_e)] \quad (2-1)$$

or by

$$r(P_e, d_e) = E_{\theta} R(\theta, d_e) = E_{\theta} [E_{z|\theta} L(z, d_e)] \quad (2-2)$$

A decision rule, d_e^* , which minimizes $r(P_e, d_e)$ is called a Bayes decision rule with respect to $P_e(\theta)$. That is,

$$d_e^* = \underset{d}{\text{Arg Min}} r(P_e, d) \quad (2-3)$$

If a decision rule, λ_e^* , minimizes $r(P_e, \lambda_e)$ subject to the restriction that

$$\underset{\theta}{\text{Max}} R(\theta, \lambda_e^*) \leq \bar{C} \quad (2-4)$$

where $\theta \in \Theta$ and \bar{C} is some specified constant, then λ_e^* is called a restricted Bayes rule with respect to P_e and \bar{C} ¹⁰. Note that Eqs. 2-3 2-4 are given for a specific experiment e .

III. GENERAL SOLUTION PROCEDURE

The crux of the problem in forming Bayes decision rules is the selection of the prior probability density for the unknown state of nature θ .

A. PRIOR PROBABILITY ASSIGNMENT

In this paper the assignment problem will be treated as an inference problem, and an objective inference technique developed by Kashyap 5, 6, 7, 8 will be primarily used.

Following this procedure, when z is the variable of interest and for finite θ , z and a given experiment e , the prior density $P_e(\theta)$ is given by

$$P_e(\theta) = \underset{P(\theta)}{\text{Arg Max}} I(z; \theta) \quad (3-1)$$

where

$$I(z; \theta) = \sum_{\theta_i} \sum_{z_j} P(z_j | \theta_i) P(\theta_i) \ln \frac{P(z_j | \theta_i)}{P(z_j)} \quad (3-2)$$

and

$$P(z_j) = \sum_{\theta_k} P(z_j | \theta_k) P(\theta_k) \quad (3-3)$$

In the maximization in Eq. 3-1, $P(\theta)$ is subject to the restrictions

$$\sum_{\theta_i} P(\theta_i) = 1 \quad (3-4)$$

$$P(\theta_i) \geq 0 \quad \forall \theta_i \in \Theta \quad (3-5)$$

and whatever testable information on θ is available

In Ref. 6 Kashyap has developed a version of Eq. 3-1 which allows for previous observations to be taken into account. If θ, z are finite and there are m independent observations from $z \rightarrow z(m+1)$ is the variable of interest now--let Z^m denote a finite valued statistic of these observations.

$$Z^m = \phi_m[z(1), z(2), \dots, z(m)] \quad (3-6)$$

Then the inferred prior for a given experiment e_m is

$$P_{e_m}(\theta) = \text{Arg Max}_{P(\theta)} I(z(m+1); \theta | Z^m) \quad (3-7)$$

where

$$I(z(m+1); \theta | Z^m) = \sum_{\theta_i} \sum_{z_k} \sum_{Z_j^m} P(\theta_i) P(z_k | \theta_i) P(Z_j^m | \theta_i) \ln \frac{P(z_k | \theta_i)}{P(z_k | Z_j^m)} \quad (3-8)$$

and

$$P(z_k | Z_j^m) = \frac{\sum_{\theta_i} P(\theta_i) P(z_k | \theta_i) P(Z_j^m | \theta_i)}{\sum_{\theta_\ell} P(Z_j^m | \theta_\ell) P(\theta_\ell)} \quad (3-9)$$

$P(z^m|\theta)$ is found from $P(z|\theta)$ and the assumption that $z(1), z(2), \dots, z(m)$ are independent. $P(\theta)$ is again subject to restrictions imposed by Eqs. 3-4 and 3-5 and the testable prior information. When $m = 0$, Eq. 3-7 just reduces to Eq. 3-1. It should be noted that if $P(z|\theta)$ is altered, the prior probabilities $P_e(\theta)$, $P_{e_m}(\theta)$ given by Eqs. 3-1 and 3-7 will be altered.

When θ is the variable of interest, the Principle of Maximum Entropy developed by Jaynes³ will be used. According to this principle, for a given experiment e , the prior $P_e(\theta)$ which maximizes the entropy of θ , $I(\theta;\theta)$ is selected as the inferred prior. For finite θ , this assignment is given by

$$P_e(\theta) = \text{Arg Max}_{P(\theta)} I(\theta;\theta) \quad (3-10)$$

where

$$I(\theta;\theta) = - \sum_{\theta_1} P(\theta_1) \ln P(\theta_1) \quad (3-11)$$

where again $P(\theta)$ is subject to Eqs. 3-4 and 3-5 and whatever restrictions imposed by testable prior information. Note the independence of $P_e(\theta)$ from $P(y|\theta)$.

A comparison of the fundamental differences between these inference techniques--Kashyap's inference and Jaynes' Maximum Entropy--is given in Ref. 7.

Equations 3-1, 3-7, and 3-10 are given for finite θ, Y . For continuous θ, Y , the summations can be replaced by integrals, and assuming these integrals exist, the same approach followed.

B. REASONABLENESS OF INFERENCE

One result of using these probability assignment procedures is that when the outcome of the experiment is also the variable of interest, the inferred density is a function of the experiment. Thus, for different experiments the densities $P_e(\theta)$ will generally differ. This is not unreasonable since the amount of information available ($P(z|\theta)$ is a form of knowledge) varies with the experiments. The implications of this will be discussed in Sec. VI.

Kashyap has justified his approach to the inference problem by showing it is the solution of a repetitive zero-sum game. Thus, if there is an unknown, but true, frequency distribution for z , any deviation from the optimal inferred density by the player can only increase his maximum possible loss. In addition, in Ref. 8 Kashyap has shown that under certain conditions his inference technique minimizes the maximum possible divergence between the unknown frequency distribution of z and that obtained using $P(z|\theta)$ and the objective prior given by Eq. 3-1. In a sense then this is a conservative inference procedure.

C. GENERAL SOLUTION

We shall use the term objective Bayes decision rules to denote those decision rules which are Bayes with respect to priors formed by using the objective inference techniques outlined in the previous section. These objective Bayes decision rules then follow directly from the definitions in Eqs. 2-3 and 2-4 and the inferred priors.

The final step in solving the general decision making problem under uncertainty is the selection of the ~~experiment~~ which results in the minimum risk for a given loss function. For the inference technique used here this necessitates comparing risks based on different priors. The validity of these comparisons is discussed in Sec. VI.

IV. EXAMPLE 1 - URN PROBLEM

This problem has been considered previously by Raiffa¹¹. We will modify it here by considering losses instead of gains and using an objective approach.

A. STATEMENT OF THE URN PROBLEM

The basic problem is as follows. There is a collection of 1000 urns, each of which is either the type θ_1 or type θ_2 , but there are no external markings distinguishing them. Urns of type θ_1 contain 4 red balls and 6 black balls; urns of type θ_2 contain 9 red balls and 1 black ball.

There are two participants---a decision maker and a neutral referee. The decision maker is to repeatedly draw one of the urns from the collection and guess whether it is of type θ_1 or θ_2 . He can also refuse to play. Depending upon his action and upon the type of urn (which is known to the referee), he receives a reward or pays a penalty. He wishes to choose his actions so that he minimizes his average loss over all the draws. After incurring the loss the decision maker returns the urn to the collection. It is assumed that the number of θ_1 , θ_2 urns in the collection does not vary so that there is an underlying (frequency) distribution for θ , albeit unknown.

Therefore, let the state space Θ_u , and action space A_u be

$$\Theta_u = \{\theta_1, \theta_2\}$$

(4-1)

$$A_u = \{a_{u1}, a_{u2}, a_{u3}\} \quad (4-2)$$

where

a_{u1} : guess θ_1

a_{u2} : guess θ_2

a_{u3} : refuse to play

The losses are given in Table 4-1.

TABLE 4-1
LOSS MATRIX FOR URN PROBLEM

	a_{u1}	a_{u2}	a_{u3}
θ_1	-40	5	0
θ_2	20	-100	0

Consider two possible experiments

e_{u0} : no observations at cost \$0.00

e_{u1} : a single observation, y , at cost \$8.00

Let y_1 = red ball; y_2 = black ball. Then

$$P(y_1|\theta_1) = .4 \quad P(y_1|\theta_2) = .9 \quad (4-3)$$

$$P(y_2|\theta_1) = .6 \quad P(y_2|\theta_2) = .1$$

where an observation y refers to withdrawing a ball from the chosen urn and observing its color.

B. SOLUTION

Following the procedure outlined in the preceding section, the objective Bayes rules $d_{e_{uo}}^*$, $d_{e_{ul}}^*$ for experiments e_{uo} , e_{ul} can easily be found.

Since the loss functions are of the form $L(\theta, d)$, the principle of Maximum Entropy will be used to infer $P_{e_{uo}}(\theta)$, $P_{e_{ul}}(\theta)$. If there are no restrictions other than

$$\sum_{i=1}^2 P(\theta_i) = 1, \quad P(\theta_i) \geq 0, \quad i=1,2 \quad (4-4)$$

then by the Maximum Entropy Principle

$$P_{e_{uo}}(\theta_1) = P_{e_{uo}}(\theta_2) = .5 \quad (4-5)$$

and

$$P_{e_{ul}}(\theta_1) = .5, \quad P_{e_{ul}}(\theta_2) = .5 \quad (4-6)$$

In order to find $d_{e_{uo}}^*$, $d_{e_{ul}}^*$, it is perhaps easiest to enumerate the candidates or allowable strategies. Denote them by σ . Then

σ_{uo} : refuse to play

σ_{u1} : choose a_{u1}

σ_{u2} : choose a_{u2}

σ_{u3} : if $y = \text{red}$, choose a_{u1}

	:	if	y	=	black,	choose	a_{u1}	
σ_{u4} :	"	"	red	"	a_{u1}			
	:	"	"	black	"	a_{u2}		
σ_{u5} :	"	"	red	"	a_{u2}			
	:	"	"	black	"	a_{u1}		
σ_{u6} :	"	"	red	"	a_{u2}			
	:	"	"	black	"	a_{u2}		(4-7)

Thus, σ_{u0} , σ_{u1} , σ_{u2} are the candidates for $d_{e_{u0}}^*$; σ_{u3} , σ_{u4} , σ_{u5} , σ_{u6} are the candidates for $d_{e_{u1}}^*$.

Following Sec. III, it is easily seen that

$$d_{e_{u0}}^* = \sigma_{u2} ; \quad d_{e_{u1}}^* = \sigma_{u5} \quad (4-8)$$

and their corresponding risks are

$$r(P_{e_{u0}}, d_{e_{u0}}^*) = .5(5) + .5(-100) = -47.5 \quad (4-9)$$

$$r(P_{e_{u1}}, d_{e_{u1}}^*) = .5(-14) + .5(-80) = -47 \quad (4-10)$$

The next step in the decision making process is the selection of the best experiment. While in this problem this selection is straightforward, it is not always so, as we shall see in the next section. Accordingly, we shall postpone selecting the "best" experiment until Sec. VI.

V. EXAMPLE 2 - BALL PROBLEM

A. STATEMENT OF THE BALL PROBLEM

The ball problem is similar to the urn problem but there are some basic differences. Again there is a collection of 1000 urns, each of which is of type θ_1 or θ_2 , having the same characteristics as in the urn problem. Again there are two participants--a decision maker and a neutral referee. However, now the decision maker is to repeatedly draw one of the urns from the collection and then withdraw a ball from this selected urn. Without knowing the type of urn, he is to guess the color of this ball. Depending upon his action and upon the color of the ball, he receives a reward or pays a penalty. Again he wishes to choose his actions so that he minimizes his average loss over all the draws. After incurring the loss, the decision maker returns the urn to the collection.

Let the state space Θ_b and action space A_b be

$$\Theta_b = \{\theta_1, \theta_2\} \quad (5-1)$$

$$A_b = \{a_{b1}, a_{b2}, a_{b3}\} \quad (5-2)$$

where

a_{b1} : guess red

a_{b2} : guess black

a_{b3} : do not play

If z denotes the observation of the ball which determines the loss incurred, then the loss function $L(z, a_b)$ is given in Table 3-2, where z_1 = red, z_2 = black.

TABLE 5-1
LOSS MATRIX FOR BALL PROBLEM

	a_{b1}	a_{b2}	a_{b3}
z_1	-40	5	0
z_2	20	-100	0

Now consider two experiments e_{b0} , e_{b1} . With e_{b0} there is only prior information on the collection of urns and the conditional probabilities $P(z|\theta)$, which are the same as those given for y in Eq. 4-3, upon which to base the decision. That is, under e_{b0} the decision maker withdraws an urn from the collection, selects an action, $a_b \in A_b$, and then withdraws a ball z from this urn. However, under e_{b1} let the decision maker withdraw a ball from the selected urn and observe its color before making a decision. Call this observation $z(1)$. Then have him replace this ball, select an action $a_b \in A_b$, withdraw another ball, $z(2) = z$, and incur the loss $L(z, a_b)$. Let the cost of this extra observation be \$3.00.

Thus, there are two possible experiments.

e_{b0} : no observations, at cost \$0.00

e_{b1} : one observation, $z(1)$, at cost \$3.00

Let $z_1 = \text{red}$, $z_2 = \text{black}$. Then $P(z_j | \theta_i)$, $i=1,2$; $j=1,2$, is the same as that given in Eq. 4-3 for y .

B. SOLUTION

Section III outlines the solution technique. $P_{e_{bo}}(\theta)$ and $P_{e_{b1}}(\theta)$ are given by Eqs. 3-1 and 3-7, respectively. If there are no restrictions other than

$$\sum_{j=1}^2 P_{e_{bj}}(\theta_i) = 1 ; \quad P_{e_{bj}}(\theta_i) \geq 0, \quad i=1,2; \quad j=1,2 \quad (5-3)$$

then

$$P_{e_{bo}}(\theta_1) = .465 ; \quad P_{e_{bo}}(\theta_2) = .535 \quad (5-4)$$

and

$$P_{e_{b1}}(\theta_1) = .505 ; \quad P_{e_{b1}}(\theta_2) = .495 \quad (5-5)$$

(Comparison with Eqs. 4-5 and 4-6 illustrates the effect of the variable of interest upon the inferred prior for θ .)

Following the same procedure as in the urn problem, we will enumerate the allowable strategies for experiments e_{bo} , e_{b1} .

σ_{bo} : refuse to play

σ_{b1} : choose a_{b1}

σ_{b2} : choose a_{b2}

σ_{b3} : if $z(1) = \text{red}$, choose a_{b1}

: " " black " a_{b1}

σ_{b4} : " " red " a_{b1}

: " " black " a_{b2}

σ_{b5} : " " red " a_{b2}

: " " black " a_{b1}

σ_{b6} : " " red " a_{b2}

: " " black " a_{b2}

Thus, σ_{bo} , σ_{b1} , σ_{b2} are the candidates for $d_{e_{bo}}^*$; σ_{b3} , σ_{b4} , σ_{b5} , σ_{b6} are the candidates for $d_{e_{b1}}^*$.

Following Sec. III, it is easily seen that

$$d_{e_{bo}}^* = \sigma_{b2}; \quad d_{e_{b1}}^*(z(1)) = \sigma_{b4} \quad (5-6)$$

and their corresponding risks are

$$r(p_{e_{bo}}, d_{e_{bo}}^*) = -58(.465) - 5.5(.535) = -29.91 \quad (5-7)$$

$$r(p_{e_{b1}}, d_{e_{b1}}^*) = -33.4(.505) - 28.15(.495) = -30.80 \quad (5-8)$$

The selection of the "best" experiment is not as straightforward as in the urn problem--since the risks in Eqs. 5-7 and 5-8 are based on different priors, how valid is a comparison between them for use in selecting the best experiment? We consider this question in the next section.

VI. SELECTION OF EXPERIMENT

The previous sections have been directed mainly at the selection of the "best" decision rule for a given experiment, where the conditional probability $P(y|\theta)$ has been fixed throughout the experiment.

However, the final step in the overall decision making problem is the selection of the optimal experiment from a set of possible experiments.

A. BASIC PROBLEM

Quite simply, the basic problem confronting us is the selection of a criterion for choosing the optimal experiment. The criterion for selection of the decision rule for a given experiment was to choose the decision rule, d_e^* , which minimized the risk $r(P_e, d_e)$ where P_e was given by our objective inference technique. Notice, however, that this risk is a function of P_e and that for loss functions of the type $L(z, a)$, P_e may depend on the experiment under consideration. Accordingly, sole use of $r(-, -)$ as a criterion--that is, choose experiment e_i over e_j if and only if

$$r(P_{e_i}, d_{e_i}^*) < r(P_{e_j}, d_{e_j}^*) \quad (6-1)$$

is difficult to justify since different probability distributions for θ are used in evaluating the respective risks. If P_{e_i}, P_{e_j} are identical, then Eq. 6-1 can easily be justified as an ordering using the arguments

presented in Part II. When these priors differ, Eq. 6-1 must be used with more care. Consider the following example.

1. Ball Problem--Choose Between e_{bo} and e_{b1}

In the previous section experiments e_{bo} , e_{b1} were introduced in the ball problem. The question now is, which experiment should be used? (Note that experiment e_{b1} has a \$3 cost while e_{bo} has no extra cost. Hence the question could be phrased as--is e_{b1} worth \$3 more than e_{bo} ?)

The respective risks for the two experiments are

$$r(P_{e_{bo}}, d_{e_{bo}}^* = \sigma_{b2}) = -58(.465) - 5.5(.535) = -29.91 \quad (6-2)$$

$$r(P_{e_{b1}}, d_{e_{b1}}^* = \sigma_{b4}) = -33.4(.505) - 28.15(.495) = -30.80 \quad (6-3)$$

However, these risks do not tell the whole story. Consider $r(p, \sigma_{b2})$ and $r(p, \sigma_{b4})$ as functions of p , the true probability that $\theta = \theta_1$.

$$r(p, \sigma_{b2}) = -5.5 - 52.5 p \quad (6-4)$$

$$r(p, \sigma_{b4}) = -28.15 - 5.25 p \quad (6-5)$$

These risks are plotted versus p in Fig. 6-1. They intersect at about $p = .480$. Thus, when $p > .480$ ($p < .480$), σ_{b2} results in less (greater) average loss than σ_{b4} , so if p actually did equal $P_{e_{b1}}(\theta_1)$ ($=.505$)

$$r(P_{e_{b1}}, \sigma_{b2}) < r(P_{e_{b1}}, \sigma_{b4}) \quad (6-6)$$

However, from Eqs. 6-2 and 6-3,

$$r(P_{e_{b0}}, \sigma_{b2}) > r(P_{e_{b1}}, \sigma_{b4}) \quad (6-7)$$

Thus, we have a situation in which we prefer σ_{b4} to σ_{b2} if we evaluate their respective risks based on $P_{e_{b1}}$, $P_{e_{b0}}$, respectively; but if we compare risks when both are evaluated at $p = P_{e_{b1}}(\theta_1)$, we prefer σ_{b2} to σ_{b4} . Obviously, factors other than just $r(P_{e_{b0}}, d_{e_{b0}}^*)$ and $r(P_{e_{b1}}, d_{e_{b1}}^*)$ should be taken into consideration.

B. QUASI-BAYES APPROACH

It seems that so long as our objective inference technique infers different priors for different experiments, our final selection of experiment must remain somewhat arbitrary. In order to formalize these arbitrary considerations somewhat, we shall introduce the quasi-Bayes approach.

We assume that there is a fixed set of experiments, $E = \{e_1, e_2, \dots, e_n\}$ from which we wish to select an experiment and decision rule. We shall call each decision rule and its associated experiment a strategy and let π_i , $i=1, 2, \dots, m$, denote the admissible strategies formed by the experiment in E and their corresponding decision rules. (Note that a decision rule d_1 might be admissible relative to the other decision rules in a given

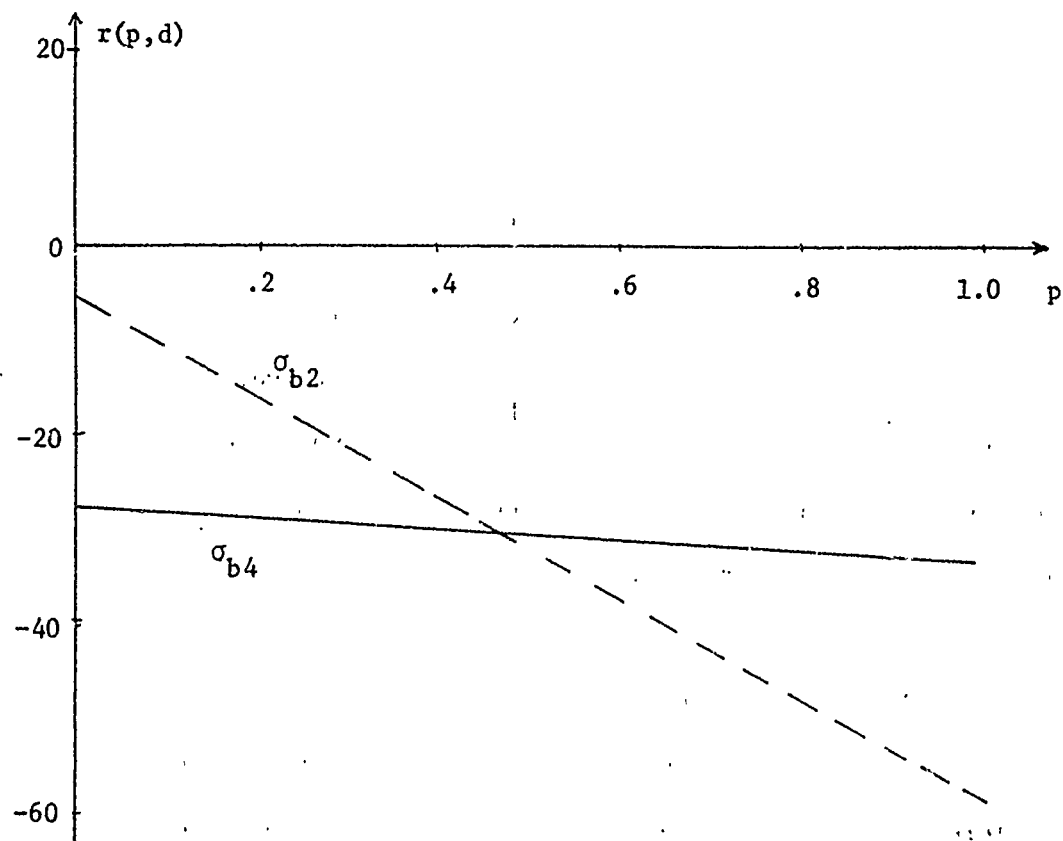


Figure 6-1: Comparison of Risks for Ball Problem

experiment $e \in E$, but the strategy (e, d_1) might be inadmissible relative to the strategies π_i , $i=1,2,\dots,m$.) Let ρ be a probability distribution over the π 's; that is, ρ is our randomized choice of decision rule and experiment. We will restrict ρ to those mixtures of π 's which are themselves admissible strategies.

Now if the true distribution for θ -- $P(\theta)$ --were known, the risk $r(P, \rho)$ could be evaluated and ρ chosen to minimize it. For example, if ρ_1, ρ_2, ρ_3 denote the probabilities that ρ selects decision rules $d_{e_1}, d_{e_2}, d_{e_3}$ from experiments e_1, e_2, e_3 , respectively, then

$$r(P, \rho) = \rho_1 r(P, d_{e_1}) + \rho_2 r(P, d_{e_2}) + \rho_3 r(P, d_{e_3}) \quad (6-8)$$

However, not only is the true distribution of θ unknown, but our objective inference technique may assign a different prior distribution, $P_e(\theta)$, to each experiment e . Using this technique, no corresponding $r(P_e, \rho)$ exists since P_e varies with the selections of ρ . Let τ denote a mapping which relates the decision rule selected using ρ with its associated objective probability. Then define $\tilde{r}(\tau, \rho)$ analogously to $r(P_e, \rho)$; note, however, that τ is not a probability density. For example, if ρ_1, ρ_2, ρ_3 are as defined in Eq. 6-8 and $d_{e_1}, d_{e_2}, d_{e_3}$ are decision rules from experiments e_1, e_2, e_3 , then

$$\tilde{r}(\tau, \rho) = \rho_1 \tau(P_{e_1}, d_{e_1}) + \rho_2 \tau(P_{e_2}, d_{e_2}) + \rho_3 \tau(P_{e_3}, d_{e_3}) \quad (6-9)$$

Define ρ^* to be a quasi-Bayes decision rule with respect to τ if

$$\rho^* = \underset{\rho}{\text{Arg Min}} \tilde{r}(\tau, \rho) \quad (6-10)$$

were naturally ρ is subject to the conditions necessary for it to be a probability distribution.

If there are no further restrictions on ρ or $\tilde{r}(\tau, \rho)$ and there is one decision rule-experiment combination whose risk is less than any other risk, then ρ will select this decision rule-experiment with probability one. That is, the unrestricted quasi-Bayes procedure just reduces to: select the decision rule d_e and experiment e having the least risk $r(P_e, d_e)$.

However, if a restriction such as

$$R(\theta, \rho^*) \leq \bar{C} \quad \forall \theta \in \Theta \quad (6-11)$$

if desired, then ρ^* will not necessarily be a degenerate probability distribution. We will denote the decision rules satisfying Eqs. 6-10 and 6-11 by $\tilde{\lambda}^*$ and call them restricted quasi-Bayes decision rules.

It is this choice of \bar{C} which allows the decision maker to formalize his arbitrary considerations. A \bar{C} close to the minimax risk indicates little confidence in the accuracy of the inferred priors, whereas a \bar{C} close to the maximum risk of the unrestricted Bayes rule indicates confidence that the inferred and actual distributions are "close."

This approach, therefore, has some of the advantages of both the minimax and unrestricted Bayes decision rules. On one hand it allows the decision maker to restrict his maximum possible loss to some level \bar{C} , but on the other hand, since P_e is used in evaluating the risk that $\tilde{\lambda}^*$ minimizes,

it isn't so pessimistic that it concentrates unduly (as determined by \bar{C}) on the worst states of nature. An example illustrating this behavior is given in the next section.

C. EXAMPLE 3 - MODIFIED BALL PROBLEM

The ball problem we will consider here has one slight modification from that presented in Secs. V and VI-A--instead of a \$3 observation cost for experiment e_{b1} , we will assume a \$6 observation cost. This modification will give a more interesting solution to the problem.

Due to this modification in experimental costs, the unrestricted quasi-Bayes rule selects experiment e_{b0} and decision rule σ_{b2} . That is, Eq. 6-2 is unchanged, but Eq. 6-3 is changed to

$$r(P_{e_{b1}}, d_{e_{b1}} = \sigma_{b4}) = -30.80 + 3 = -27.80 \quad (6-12)$$

and consequently $r(P_{e_{b0}}, \sigma_{b2})$ is less than this new risk so (e_{b0}, σ_{b2}) is the unrestricted quasi-Bayes rule.

If we wish to restrict the maximum possible risk to less than some specified constant \bar{C} , it is then necessary to randomize over experiments. Of the set of strategies available, only σ_{b1} , σ_{b2} , and σ_{b4} are admissible. If $\tilde{\lambda}_1^*$, $\tilde{\lambda}_2^*$, $\tilde{\lambda}_3^*$ represent the probabilities that the restricted quasi-Bayes rule $\tilde{\lambda}^*$ selects σ_{b1} , σ_{b2} , and σ_{b4} , respectively (and their associated experiments), then applying Eqs. 6-10 and 6-11--for this case a linear programming problem results--we get

$$\tilde{\lambda}_1^* = 0, \quad \tilde{\lambda}_2^* = \frac{\bar{C} + 25.15}{19.65} = .263, \quad \tilde{\lambda}_3^* = .737 \quad (6-13)$$

if $\bar{C} = -20$; if $\bar{C} = -25.15$, we get

$$\tilde{\lambda}_1^* = 0 = \tilde{\lambda}_2^*, \quad \tilde{\lambda}_3^* = 1 \quad (6-14)$$

and if $\bar{C} = -26$, we get

$$\tilde{\lambda}_1^* = -\frac{\bar{C} + 25.15}{8.85} = .096, \quad \tilde{\lambda}_2^* = 0, \quad \tilde{\lambda}_3^* = .904 \quad (6-15)$$

Note that the range for effective \bar{C} is

$$-26.43 \leq \bar{C} \leq -5.5 \quad (6-16)$$

That is, \bar{C} should be greater than the minimax risk -26.43 and less than the maximum risk under the unrestricted quasi-Bayes rule.

This minimax risk can be obtained by using well established techniques. If γ^* denotes the minimax strategy and $\gamma_1^*, \gamma_2^*, \gamma_3^*$ the probabilities that γ^* respectively selects $\sigma_{b1}, \sigma_{b2}, \sigma_{b4}$, (and, of course, their corresponding experiments) then applying these techniques we get

$$\gamma_1^* = .149, \quad \gamma_2^* = 0, \quad \gamma_3^* = .851 \quad (6-17)$$

and the corresponding minimax risk is -26.43 , independent of the true probability distribution of θ since γ^* is an equalizer rule.

For these rules the risk $r(p, d)$ can be calculated where p is the actual distribution of θ . In Fig. 6-2 the risks for these rules are plotted versus p , the true probability that $\theta = \theta_1$. From the figure it is seen that the decision maker has choices ranging from a minimum possible risk of $-\$58$, with a possible risk of $-\$5.5$, to a guaranteed risk of $-\$26.43$. The compromising behavior of the restricted quasi-Bayes is evident.

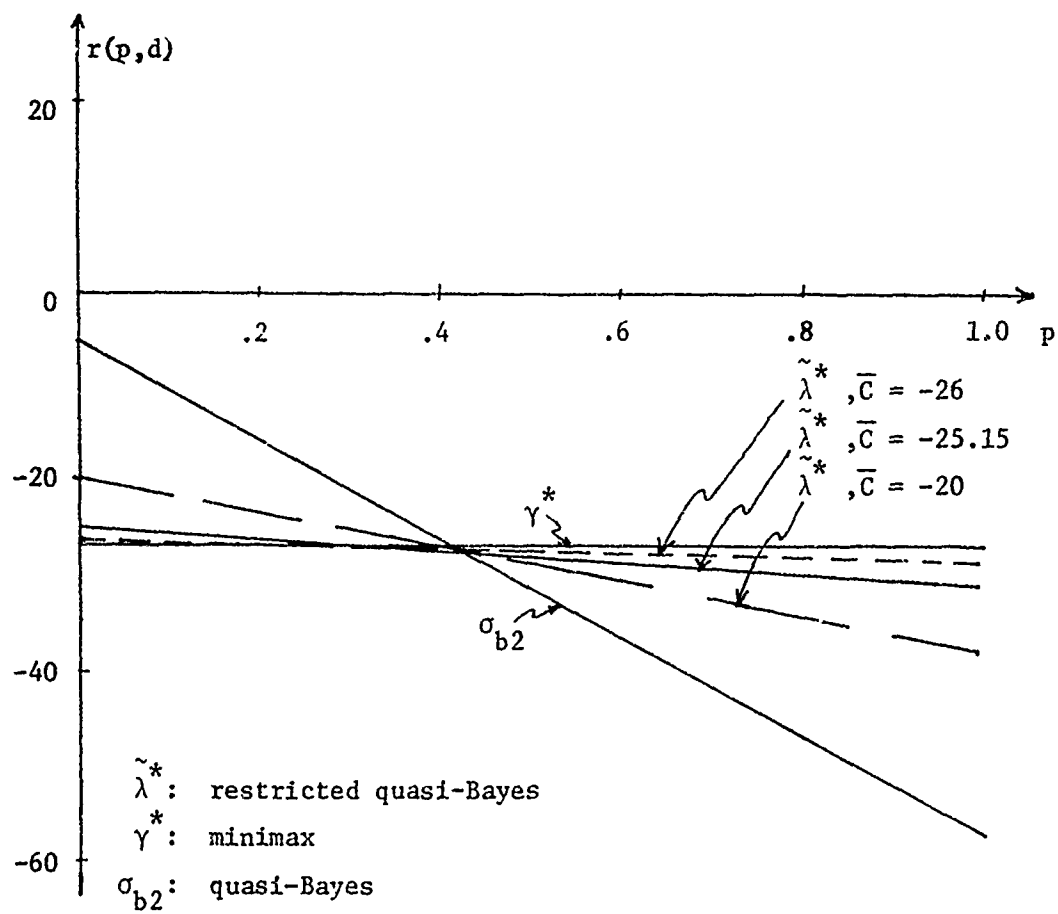


Figure 6-2: Comparison of Risks for Ball Problem with \$6 Observation Cost

D. DISCUSSION

The basic difficulty, of course, is the possible variation in inferred prior with changing $P(y|\theta)$. If the inferred prior, $P(\theta)$, were fixed over the experiments, $r(P,d)$ could be used as the sole decision criterion. However, a variation in the prior requires some modification in the decision making procedure. (We discuss some of the anomalies associated with quasi-Bayes in Part II of this paper.) We have elected to use the \bar{C} restriction, feeling that it is a concise and general way of handling the difficulty. In general any Bayes procedure using an (objective) inference technique which in turn is a function of $P(y|\theta)$ will encounter this difficulty of multiple priors.

VII. CONCLUSIONS

The use of this approach enables the solution of a wide range of practical problems involving uncertainty. It can be applied to the group Bayesian problem (Raiffa , Chapter 8)--collectively selecting \bar{C} may be easier than finding a collective (subjective) prior. In Ref. 12 this procedure is used to find both open-loop and closed-loop controllers for dynamic systems having unknown parameters. An advantage of this approach is its ability to objectively use the available prior information.

When there is only one experiment, the application of our approach is straightforward and, as will be seen in Part II, quite justifiable. It is, in a sense, an objective alternative to the minimax procedure. However, when there is a choice of experiments, our results are not as conclusive. More will be said on this subject in Part II but suffice it to say that no alternative decision making procedure exists which doesn't have some serious objections to it!

REFERENCES

1. L. J. Savage, The Foundation of Statistics, New York: Wiley, 1954.
2. H. Jeffreys, Theory of Probability, 2nd ed., Oxford: Clarendon Press, 1948.
3. E. T. Jaynes, "Prior Probabilities", IEEE Trans. Syst. Sci. Sybern, Vol. SSC-4, September 1968, pp. 227-241.
4. M. Tribus, (1969); "Rational Descriptions, Decisions and Designs", Pergamon Press.
5. R. L. Kashyap, "Prior Probability, Two-Person Games and Information Theory", TR-EE 69-44, School of Electrical Engineering, Purdue University, 1969.
6. -----, "A Theory of Statistical Inference Based on a Stochastic, Two-Person, Zero-Sum Game", unpublished manuscript.
7. -----, "Prior Probability of a Discrete Variable Using Two-Person Game Theory", unpublished manuscript.
8. -----, "Prior Probability and Uncertainty", IEEE Trans. Inform. Theory, Vol. IT-17, Nov. 1971, pp. 105-113.
9. T. S. Ferguson, Mathematical Statistics: A Decision Theoretic Approach, New York: Academic Press, 1967.
10. J. L. Hodges, Jr. and E. L. Lehmann, "The Use of Previous Experience in Reaching Statistical Decisions", Ann. Math. Stat., Vol. 23, 1952, pp. 396-407.
11. H. Raiffa, Decision Analysis, Reading, Mass.: Addison-Wesley, 1968.
12. T. Oberlin, Bayes Decision Rules Based on Objective Priors, Chapter 7, Ph.D. thesis, Purdue University, January 1972.

BAYES DECISION RULES BASED ON OBJECTIVE PRIORS

PART II: JUSTIFICATION

By T. L. Oberlin and R. L. Kashyap

Abstract

In Part I of this paper an objective Bayes approach to decision making under uncertainty was proposed. The priors were obtained by using an objective inference technique developed by Kashyap. Justification arguments are now given for this approach based upon axiomatic considerations, reasonableness arguments and comparison of the average losses incurred using both this approach and the minimax approach.

It is shown that in situations where the experimentation is fixed and the decision problem is faced repeatedly, but not necessarily an infinite number of times, this approach is justifiable. In situations where there is a choice of experiments, these arguments are not as conclusive; however, the approach still has practical merit as an objective alternative to the minimax approach.

I. INTRODUCTION

The basic approach of Part I was to assume that the uncertainties present in the decision problem could be characterized by random variables which are then assigned probability distributions.

We wish now to give justification for this approach. We shall do this by axiomatic considerations, reasonableness arguments and comparison of average losses incurred using other approaches. These approaches include subjective Bayes, empirical Bayes (Robbins¹) and minimax. Techniques such as hypothesis testing, confidence intervals and point estimation will not be discussed. Lindley² gives an interesting comparison between some of these approaches and those using the concept of a prior distribution.

The merits of and objections to each approach depend a great deal upon the type of problem in which they are used, especially the repetitive nature of the problem. We start by considering a repetitive decision problem in which there is no choice of experiment and then generalize to the situation where the decision maker has a choice of experiments.

II. SINGLE EXPERIMENT

In general we do not expect to show that any one decision rule is always "better" than others; that is, there will be situations in which a number of decision rules are admissible. Naturally, we are primarily interested in those situations in which the objective Bayes rules seem to be better than any others, so we assume that we are in a situation in

which the decision problem is faced repeatedly--not necessarily an infinite number of times--and in which the experimentation is fixed.

In order to provide a basis for comparison and to provide insight into the nature of the justification problem, an axiomatic approach to selecting a decision criterion is briefly discussed, and then arguments are given for the validity of other approaches to forming decision rules. Finally, a direct comparison of the objective Bayes and minimax approaches is made.

A. AXIOMATIC APPROACH

In using the axiomatic approach, the order of the decision-making process is reversed from that given in Part I. In Part I a decision criterion is postulated, resulting decision rules formed and then the performance of these decision rules is evaluated. An alternative procedure--the axiomatic approach--is to list certain desiderata or axioms for a decision criterion to fulfill and then see if the proposed criterion satisfies these axioms. This is done prior to using the criterion to find decision rules. The axioms are posed, naturally, so that if a criterion fulfills them the resulting decision rules are justifiable. Not surprisingly, one of the major problems in applying the approach is the selection of axioms such that they are mutually compatible and also intuitively reasonable (see Birnbaum³).

In Chapter 13 of their book⁴, Luce and Raiffa give a very readable presentation of the axiomatic approach, and consequently it will not be

repeated here. However, some conclusions can be drawn. Namely, of the four decision criteria that they applied to their list of axioms--minimax, minimax regret, α pessimism-optimism and insufficient reason--no one criterion dominated. That is, of these four decision criteria, none fulfilled all the axioms in a completely satisfactory manner.

However, the criterion based on the principle of insufficient reason satisfied the basic axioms 1 thru 9 in Luce and Raiffa, but there is question as to when the criterion is applicable, that is, the axiomatic representation of complete ignorance. The essence of the argument is that if the decision maker is "completely ignorant" as to the state of nature, then should his final decision rule be altered by deleting a state of nature (assuming Θ finite) that has the same losses as another state for any action $a \in A$? If the number of states of nature is changed, the probability assigned to each state by the principle of insufficient reason will change, which can change the selected decision rule. Hence, the argument is joined!

If the decision criterion given by the objective Bayes approach is applied to these axioms, for fixed experimentation (only one prior $P_e(\theta)$) it can be shown that it satisfies the basic axioms 1 thru 9. However, it is subject to the same criticism as the insufficient reason criterion. Namely, that if the number of states of nature is changed, the inferred prior $P_e(\theta)$ will generally change and possibly the resulting objective Bayes decision rule will be altered.

It seems as though any decision criterion based upon an inference approach which is independent of the values assigned to the losses--as both the principle of insufficient reason and our objective inference are--will be subject to this criticism. For these inference techniques, the number of states of nature must remain fixed for a given problem; if this number changes, the problem is changed and a new ordering of decision rules is possible. It is felt that this failing is relatively minor in comparison with the difficulties associated with those procedures--such as minimax or minimax regret--which can be viewed as having inference techniques based upon the loss function.

In summary, the aniconic approach is useful since it shows that one has yet found a decision criterion which has all the desired properties that can logically be desired.

B. JUSTIFICATION OF BAYES APPROACH

If a Bayes approach is to be followed--that is, θ is assumed to be a random variable and a probability distribution is assigned to it--two basic questions concerning justification arise⁵:

- (i) adequacy of the assumption of randomness of θ and knowledge of the prior probability distribution of θ ;
- (ii) interpretation of the minimum expected loss as optimum.

Obviously the answers to (i), (ii) will depend on the probability distribution assignment technique. The following sections will discuss several of these assignment techniques, including Kashyap's, in light of (i), (ii).

1. Subjective Bayes

Essentially a subjectivist holds the view that probability measures a person's degree of belief as evidenced by his betting or action behavior. DeFinetti⁶ demonstrates that if a person is consistent in placing his bets, his subjective probability assignments will satisfy the usual laws of probability. From this point of view then, (i) is adequately answered by the subjective assignment of a probability distribution to θ .

In order to answer question (ii), the subjectivist views the loss function as a negative utility in the sense of Von Neumann-Morgenstern. In Ref. 7 Savage showed these concepts of utility and subjective probability can be linked together. Now the optimality of the minimum expected loss, where the expectation is with respect to the subjective probability distribution, follows from the expected utility hypothesis. Thus, question (ii) is answered.

There are some objections to this approach. One is the practical difficulty in assigning the probability distribution such that there are no inconsistencies. A more fundamental objection is that two different people, when faced with the same decision problem, may end up with two completely different solutions.

All in all, however, in a non-repetitive situation--that is, the decision problem is a "one shot" affair--the subjective approach is hard to argue against, even considering the difficulties of consistent assignment and nonobjectivity. However, in repetitive type situations where

there is a true underlying frequency probability distribution for θ --albeit unknown--the subjective approach seems hard to defend. It has no guarantee against the improper use of information--the actual average incurred loss may differ widely from the expected loss. Hence, the argument for answering (ii), while valid for non-repetitive situations, is not satisfactory for repetitive situations.

2. Empirical Bayes

The empirical Bayes approach is suitable for situations in which the decision problem presents itself repeatedly and independently with a fixed but unknown underlying frequency distribution for θ . A good summary of the method is given by Robbins¹. (Maritz⁸ also summarizes the method and contains more current references.) Essentially, in using the empirical Bayes approach, the decision maker forms sequential decision rules which utilize past observations to extract some information about the frequency distribution of θ . Under certain conditions then it can be shown that as $n \rightarrow \infty$ (n is the number of repetitions), the risk corresponding to these decision rules converges to the actual Bayes risk that results if the frequency distribution of θ were known. (This property is called asymptotic optimality of the decision rule.)

Consequently, in the limit as $n \rightarrow \infty$, (i) and (ii) can be answered by the same arguments given for when the frequency distribution of θ is known; that is, by the law of large numbers, the Bayes risk approaches the actual average loss.

The major objection to this approach is that the value of n necessary for convergence is not clearly specified. In practical situations there may be only a limited number of observations available. In these cases, general results for the assessment of the performance of empirical Bayes methods have not yet been found (p. viii, Ref. 8).

3. Objective Bayes

In using the objective Bayes approach, we assume that the problem is repeated often enough so that a purely subjective approach is unsatisfactory. In general, the situation can be characterized as one in which the decision problem presents itself repeatedly and independently with the same underlying frequency distribution of θ , but the number of repetitions is insufficient for convergence of an adaptive decision procedure such as empirical Bayes. For a problem of this type, initial choices for the unknown probability distributions will have a pronounced effect on the average loss incurred.

In a sense then, the question of randomness of θ is answered by the type of problem considered-- θ is assumed to have an unknown frequency probability distribution. The answer to the second part of question (1) lies in the inference technique used. Since the actual performance of the decision rule can be noticeably affected by the inferred prior probability distribution of θ , a conservative approach is desirable. Kashyap's inference technique has this conservative character. Thus, since the true frequency distribution of θ is unknown, a conservative "estimate" of it is used instead.

Note that this estimate is conservative from the inference point of view and is not necessarily conservative as far as the decision problem is concerned, that is, the actual average loss incurred. This leads to question (ii)--what is the interpretation of the resulting risk $r(P_e, d_e^*)$? $r(P_e, d_e^*)$ is really just an estimate of the average loss that will be incurred using d_e^* , but due to the nature of the problem its accuracy cannot be verified, just placed within limits. Some of the difficulties in using $r(P_e, d_e^*)$ as the sole measure of performance are discussed in Sec. III.

4. Restricted Objective Bayes

If the decision maker is not satisfied with the bounds--that is, the maximum possible average loss is too high--he may wish to choose a decision rule γ^* to minimize his upper bound, that is, his maximum possible average loss. Of course, γ^* is simply a minimax decision rule. However, this minimax approach has certain drawbacks.

A compromise between these two approaches--objective Bayes and minimax--is offered by the restricted objective Bayes approach. This approach has some of the advantages of both the minimax and unrestricted Bayes decision rules. On one hand it allows the decision maker to restrict his maximum possible loss to some level \bar{C} , but on the other hand, since P_e is used in evaluating the risk that λ_e^* minimizes, it isn't so pessimistic that it concentrates unduly (as determined by \bar{C}) on the worst states of nature. Examples illustrating this behavior are given in Part I.

In a sense then, our restricted objective Bayes rule is a hybrid-- it is based on an objective prior probability and loss function but is affected by an arbitrary value of \bar{C} . However, since this arbitrariness is restricted to the choice of \bar{C} at the "end" of the problem solution rather than in the beginning as is the case with the subjective approach, this approach avoids some of the criticisms of the purely subjective approach. Namely, it is applicable to situations in which the decision problem is faced repeatedly, and the effect of the subjectivity is easily seen. Ultimately the decision problem is subjective, and this seems to be a good place to account for it.

Undoubtedly there are those who will disagree with some of the arguments and assumptions used above. In the next section the actual performances of objective Bayes rules and minimax rules will be compared to provide a more real or operational justification of our approach.

C. COMPARISON WITH MINIMAX DECISION RULES

For the type of problem being considered--repetitive but not necessarily having an infinite number of repetitions of the decision problem-- it is felt that the minimax decision rule is the main competitor of the objective Bayes decision rules.

1. Basic Differences

The minimax approach is the most conservative decision making approach--it can be given a zero-sum game interpretation.

A comparison with the objective Bayes decision rules can be made by viewing these minimax decision rules as simply Bayes rules with respect to the least favorable distribution over θ . The difference between the two approaches then lies in the prior distributions of θ that are assumed. In the objective Bayes approach, Kashyap's inference technique is used, which itself is an interpretation of the solution of a two-person zero-sum game. However, the essential difference is that this inference procedure is independent of the values in the loss function; that is, $P_e(\theta)$ does not depend on the losses assigned in the decision problem. However, as noted it does depend on the form of the loss function. This is not the case for the least favorable distribution, which obviously depends upon the loss function values.

It is felt that the major objection to minimax decision rules--their high sensitivity to greater risk with low sensitivity to less risk--arises because of this dependence of the inference of the loss function. By avoiding this dependence, the objective Bayes decision rules are not as subject to this criticism. The following example will illustrate this.

Consider the urn problem described in Part I and the single experiment e_{u1} .

If γ^* denotes the minimax decision rule for this problem, where γ^* selects σ_{u3} and σ_{u5} with probabilities γ_1^* , γ_2^* , respectively, then

$$\gamma_1^* = .524$$

$$\gamma_2^* = .476$$

(2-1)

The corresponding minimax risk is

$$r(P^W, \gamma^*) = -23.43 \quad (2-2)$$

The objective Bayes decision rule for experiment e_{ul} is

$$d_{e_{ul}}^* = \sigma_{u5} \quad (2-3)$$

with corresponding risk

$$r(P_{e_{ul}}, d_{e_{ul}}^*) = -47 \quad (2-4)$$

Comparison of $r(P^W, \gamma^*)$ and $r(P_{e_{ul}}, d_{e_{ul}}^*)$ is not too meaningful since these risks are evaluated at different probability distributions for θ . A more meaningful comparison is to consider their respective risks evaluated with the true probability of $\theta_{1,p}$. These risks can easily be evaluated.

$$\begin{aligned} r(p, d_{e_{ul}}^*) &= 66p - 80 \\ r(p, \gamma^*) &= -23.43 \end{aligned} \quad (2-5)$$

They are plotted versus p in Fig. 2-1.

Notice that if $p < .857$, the risk incurred using the objective Bayes rule is less than that incurred using the minimax rule. In fact, the objective Bayes rule can be almost four times less than the minimax risk (-80 versus -23.43 when $\theta = \theta_1$). On the other hand, the objective Bayes risk can go as high as -14 versus the -23.43 bound on the minimax risk.

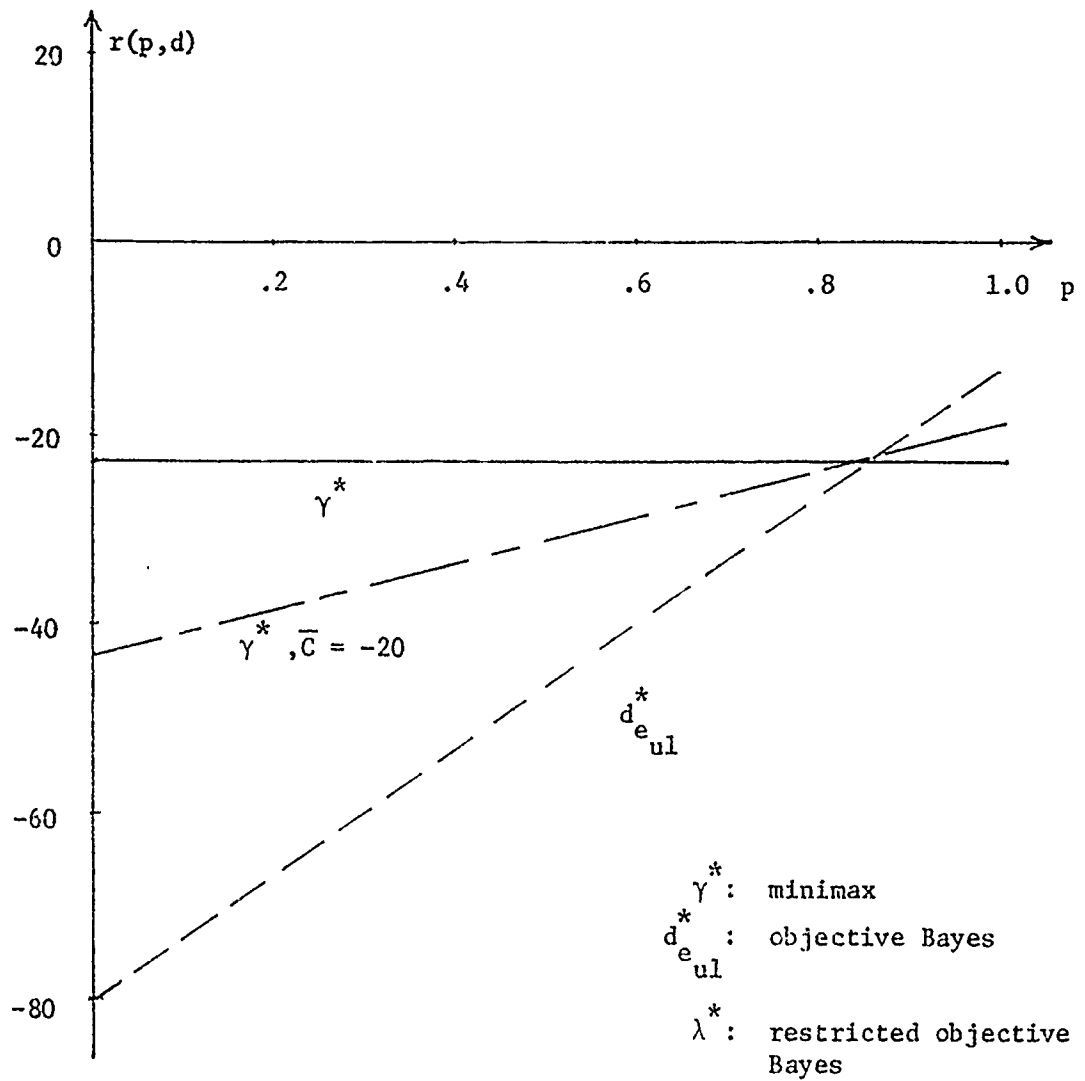


Figure 2-1: Comparison of Risks for Urn Problem

This then demonstrates the sensitivity of the minimax approach-- highly sensitive to greater risks, much less sensitive to lesser risks. This is the major criticism of the minimax approach.

However, if the decision maker objects to the possibility of an average loss of -\$14, he can use a restricted objective Bayes rule to reduce this maximum risk. Let \bar{C} denote the desired maximum possible risk. Then the restricted objective Bayes rule λ^* for experiment e_{u1} is

$$\lambda_1^* = -\frac{\bar{C} + 14}{18}, \quad \lambda_2^* = \frac{32 + \bar{C}}{18} \quad (2-6)$$

where $-23.43 \leq \bar{C} \leq -14$ and λ_1^*, λ_2^* represent the probabilities that λ^* selects σ_{u3} and σ_{u5} respectively.

For $\bar{C} = -20$, the risk $r(p, \lambda^*)$ is given by

$$r(p, \lambda^*) = 24p - 44 \quad (2-7)$$

This is plotted in Fig. 2-1. Notice that in exchange for a smaller maximum risk, the minimum possible risk is reduced. As \bar{C} approaches -23.43, these upper and lower bounds converge and the restricted objective Bayes rule approaches the minimax rule. Thus, the restricted objective Bayes rule can be viewed as a sort of compromise between the objective Bayes and minimax decision rules.

Another consideration for use in the comparison of objective Bayes and minimax decision rules is the use of sampling information. There are various ways for the objective Bayes approach to use this information, essentially by the use of judicious approximations.

However, a true minimaxer can't really make even these approximations: he is only concerned with the possibility that $P(\theta)$ can assume its least favorable value and not the probability that it does.

Rådner and Marshak⁹ give a good example illustrating some of the objections to using minimax rules when sampling is allowed. They essentially have found a situation in which even though $P(\theta)$ is unknown, the optimum minimax decision rule calls for no sampling, no matter how cheaply it may be obtained. The problem is that as long as it is possible for $P(\theta)$ to be least favorable, the minimax approach will assume it is so.

Thus, in summary, the attitude of the minimax approach makes it difficult to use "uncertain" information.

2. Variations in the Loss Function

So far we have assumed that the loss function is a known, scalar performance index. However, if we allow unknown, random terms in the loss function, the sensitivity of the minimax approach can be considerably increased relative to the objective Bayes approach.

Another shortcoming in the minimax approach arises when multiple losses are considered; that is, the performance index may be given by

$$PI = E \left(\begin{array}{c} R_1(\theta, d_{1,n}) \\ R_2(\theta, d_{1,n}) \\ \vdots \\ R_m(\theta, d_{1,n}) \end{array} \right) \quad (2-8)$$

or

$$PI = \Pr [R_1(\theta, d_{1,n}) \geq \alpha_1 \cap R_2(\theta, d_{1,n}) \geq \alpha_2 \cap \dots \cap R_m(\theta, d_{1,n}) \geq \alpha_m] \quad (2-9)$$

where $d_{1,n}$ represents a collection of decision rules (d_1, d_2, \dots, d_n) .

The objective Bayes approach is suitable for application to decision problems having performance indices of this form due to the independence of the inference from the loss function. The minimax approach, on the other hand, isn't really suitable for these types of performance indices since evaluation of either Eqs. 2-8 or 2-9 requires a probability distribution for θ . However, for these performance indices, the minimax approach doesn't indicate a way to find this probability distribution.

D. SUMMARY

This section has been concerned with the justification, for field experimentation, of the objective Bayes approach. Part B contained reasonableness arguments--what properties should a decision criterion possess;

what interpretation can be given to the resulting expected losses? In this section a more practical approach was taken--is there another objective approach that does as well or better than the objective Bayes approach?

Based on the answers to these questions and the number of situations in which the objective Bayes approach is better than the minimax approach, we feel the objective Bayes is a justifiable approach for fixed experimentation. In the next section the question of its justification when there are several priors (that is, several different experiments) in the decision problem is considered.

III. MULTIPLE EXPERIMENTS

Assume that there are a finite number of experiments from which to choose. The basic difficulty confronting us is the possible variation in priors with experiment. ($P(y|\theta)$ is fixed throughout an experiment, but may vary between experiments and consequently so can the objectively inferred prior.)

To attack this problem, we formulated the quasi-Bayes approach. We are now interested in the validity of this approach. We shall use much the same procedure as in Sec. II--axiomatic considerations, reasonableness arguments and direct comparison with the minimax approach. However, in order to illustrate some of the properties of the quasi-Bayes approach, we shall first introduce a new approach--the standard approach--and compare the two.

A. ALTERNATE APPROACH TO SELECTION OF EXPERIMENT

1. Standard Approach

A typical way of selecting an experiment and ultimately the "optimal" decision rule is to find the best decision rule for each experiment and then select the experiment whose associated best decision rule has the minimum risk. We shall call this approach the standard approach.

We can formalize this approach as follows. For each experiment $e_i \in E$, $i=1,2,\dots,n$, let the associated objective Bayes decision rule be denoted by $\lambda_{e_i}^*$. That is,

$$\lambda_{e_i}^* = \text{Arg Min}_{\lambda_{e_i}} r(P_{e_i}, \lambda_{e_i}) \quad (3-1)$$

where P_{e_i} is the objective prior probability density for experiment e_i found using Kashyap's inference technique.

In addition there may be restrictions of the form

$$R(\theta, \lambda_{e_i}^*) \leq \bar{C}_S, \quad \forall \theta \in \Theta, \quad i=1,2,\dots,n \quad (3-2)$$

on these decision rules.

Now let δ be a randomized rule for selecting the experiment, where δ_i , $i=1,2,\dots,n$, correspond to the probability that experiment e_i is selected. If

$$\hat{r}(\delta) = \delta_1 r(P_{e_1}, \lambda_{e_1}^*) + \dots + \delta_n r(P_{e_n}, \lambda_{e_n}^*)$$

then following the standard approach, the rule for selecting the experiment is given by δ^* where

$$\delta^* = \underset{\delta}{\text{Arg Min}} \hat{r}(\delta) \quad (3-3)$$

(We have introduced this notation to facilitate comparison with the quasi-Bayes approach. Obviously δ^* will be a degenerate probability distribution which simply selects the experiment having minimum $r(P_e, \lambda_e^*)$.)

2. Comparison of Approaches

The basic difference between the two approaches is that with the standard approach an experiment is selected only on the performance of its "best" decision rule whereas with the quasi-Bayes approach an experiment can be selected on the performance of any decision rule leading to an admissible strategy. (A strategy is any decision rule-experiment combination.) There are two major consequences of this difference that are of interest to us: the admissibility of the experiment and consequent decision rule which are chosen, and the difference in ranges of \bar{C}_S and \bar{C} .

If the inferred prior $P_e(\theta)$ is the same for all $e \in E$, then the experiment and consequent decision rule selected by the standard approach will be admissible relative to the other possible experiment-decision rule combinations. However, if the inferred prior $P_e(\theta)$ varies with the experiment, then it is quite possible that the experiment-decision rule selected by the standard approach will be inadmissible relative to the other strategies. Using the quasi-Bayes approach, it is easy to avoid

this by initially considering only admissible strategies; that is, we restrict ρ to admissible strategies and admissible mixtures of these strategies. Consequently, the resulting ρ^* is admissible.

Another difference between the approaches is the ranges of \bar{C} , \bar{C}_S . Specifically, it is possible for

$$\min(\bar{C}) < \min(\bar{C}_S) \quad (3-4)$$

where $\min(-)$ means that if \bar{C} or \bar{C}_S is less than this value, no solutions exist which satisfy the constraint.

It is easy to see this from a graphical point of view. Assume $E = \{e_0, e_1\}$ and that there are three admissible strategies for each experiment-- σ_{0i}, σ_{1i} , $i=1,2,3$ --having conditional risks $R(\theta, \sigma)$ as plotted in Fig. 3-1. Only the lower left hand boundaries of the risk sets are drawn.

From Fig. 3-1 we see that if $\bar{C} = \underline{C}$, we can obtain a minimax risk of \underline{C} using the restricted quasi-Bayes approach with $\bar{C} = \underline{C}$. However, with the standard approach, if $\bar{C}_S < \underline{C} + \Delta$, Fig. 3-1 shows that no solutions $\lambda_{e_0}^*, \lambda_{e_1}^*$ exist which satisfy constraint (3-2).

Moreover, using the quasi-Bayes approach with

$$\underline{C} < \bar{C} < \underline{C} + \Delta \quad (3-5)$$

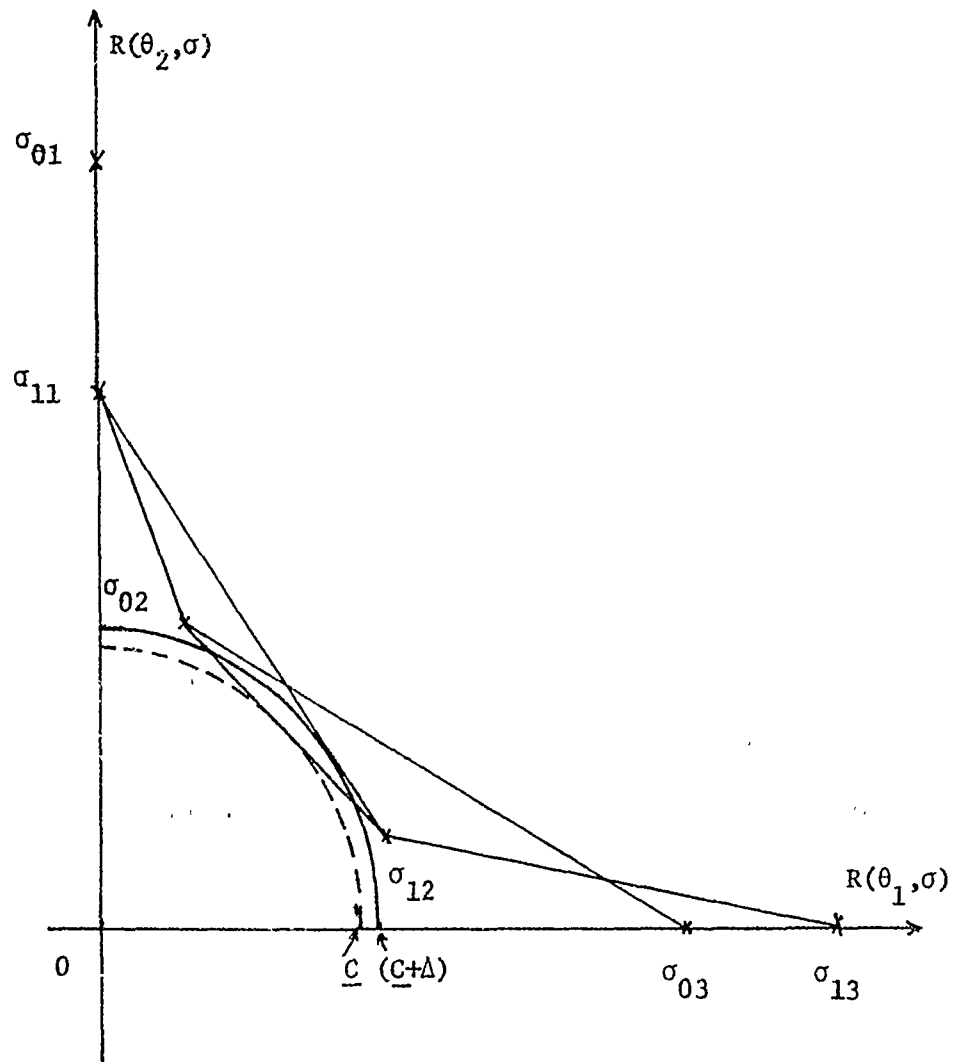


Figure 3-1: Conditional Risks for e_0, e_1 Strategies

we can obtain restricted quasi-Bayes solutions whereas no solutions exist for corresponding values of \bar{C}_S using the standard approach. (Note that this is true regardless of the priors assigned to θ .)

If \bar{C} , \bar{C}_S are so large as to be ineffective, then $\lambda_{e_i}^*$, $i=1,2,\dots,n$, will be nonrandom decision rules and the standard approach will yield the decision rule and experiment having minimum $r(P_{e_i}, \lambda_{e_i}^*)$. Assuming that this experiment-decision rule is an admissible strategy, the quasi-Bayes approach will yield the same decision rule and experiment.

Now if \bar{C} , \bar{C}_S are of such values that they are effective, it is hard to compare the resulting selections. The standard approach will select an experiment and then randomize over possible decision rules for this experiment. The quasi-Bayes approach will randomize over both experiment and decision rule; in fact, when a strategy is selected, the decision rule to be used in the experiment is specified. Consequently, following the quasi-Bayes approach, an experiment may be selected and performed and a decision rule used which might not be optimal if that experiment alone were being considered. Thus, if \bar{C} , \bar{C}_S are such that they are effective, the standard approach will yield decision rules which are consistent within the chosen experiment, and the quasi-Bayes approach can yield decision rules which, while not necessarily the best for the given experiment, are optimal, in the sense defined, over the set of experiments.

Note that when we say the decision problem is faced repeatedly we mean that the selection of both the experiment and the decision rule is faced repeatedly. Consequently, we feel that the quasi-Bayes approach is a "better" approach than the standard approach.

B. AXIOMATIC CONSIDERATIONS

The axiomatic approach was briefly discussed in Sec. II. It was stated there that the axioms presented by Luce and Raiffa were satisfied by the objective Bayes criterion, with the exception that the ordering is altered if the number of states is changed by deleting a state having the same losses as another state for any action $a \in A$. Unfortunately, the extension of the objective Bayes--the quasi-Bayes approach--does not satisfy all the basic axioms.

For example, admissibility of the optimal decision rule is a basic axiom, and we can only obtain it by restricting ρ to admissible strategies and admissible mixtures of strategies. Without this restriction the final choice of experiment-decision rule might not be admissible.

This restriction can cause some other discrepancies, however. For example, if two strategies are optimal, that is, quasi-Bayes decision rule then a mixture ρ of them may have the minimum $\tilde{r}(\tau, \rho)$, but this mixture may be inadmissible and hence non-optimal. In other words, the optimal set of strategies is not necessarily convex.

There are other discrepancies. If we add a constant to a row of the matrix (assuming finite Θ, A), we can change the ordering between strategies. For example, suppose \$100 is added to the second row of the loss matrix for the ball problem, Table 5-1, Part I. Then the priors $P_{e_{bo}}$, $P_{e_{b1}}$ remain unchanged but the risks $r(P_{e_{bo}}, \sigma_{b2})$, $r(P_{e_{b1}}, \sigma_{b4})$ go from -29.91 and -30.80 to 3.34 and 4.50, respectively. Hence, the (unrestricted) quasi-Bayes rule is changed from σ_{b2} to σ_{b4} by adding a constant to a row of the loss matrix.

While the failure of the quasi-Bayes approach to satisfy these axioms is discouraging, it should not preclude the use of this approach, since, as mentioned previously, no one decision criterion has been found which has all the desired properties as expressed by the axioms. In view of this, the ultimate justification of this approach is how well it performs relative to the alternative approaches. (It should be noted that the standard approach exhibits these same faults.)

C. REASONABLENESS ARGUMENTS

In Sec. II arguments were presented for the reasonableness of certain Bayes procedures for different situations. Similar arguments can be given for using the quasi-Bayes approach to select the best experiment. The repetitive nature of the problem is the same (except now, of course, the decision maker repetitively selects his experiment) and so conservative estimates of the prior probability of θ for each experiment are desirable. Again, the interpretation of the resulting risk is arbitrary; the risk can be bounded, but otherwise no properties such as convergence to the actual average loss can be claimed.

In Part I, Sec. VI, an example was solved using both quasi-Bayes and minimax decision rules. From Fig. 6-2, Part I, it can be seen then that the quasi-Bayes approach is a viable one (at least for this particular problem). A comparison between the two approaches was made for another problem--control of a dynamic system with unknown parameters¹⁰--with much the same type results.

In general it is felt that often the quasi-Bayes approach will yield experiment-decision rule choices whose performance is quite satisfactory. From a practical point of view, this fact in itself is sufficient to justify its use.

IV. CONCLUSIONS

Well now, where do we stand with respect to justification of our approach to decision making under uncertainty? We have looked at alternate approaches in terms of axiomatic considerations, reasonableness arguments and comparison of average losses incurred. The conclusion that we now reach is that the objective Bayes (and its generalization the quasi-Bayes) approach is a good practical approach to this problem.

There are inconsistencies in this approach; however, the majority of them only occur when the inferred prior varies with the experiment. If the inferred prior is the same for all experiments under consideration, the justification arguments presented in Sec. I for the case of a single experiment apply to the quasi-Bayes approach.

In Part I, Sec. VI, an example was solved using both quasi-Bayes and minimax decision rules. From Fig. 6-2, Part I, it can be seen then that the quasi-Bayes approach is a viable one (at least for this particular problem). A comparison between the two approaches was made for another problem--control of a dynamic system with unknown parameters¹⁰--with much the same type results.

In general it is felt that often the quasi-Bayes approach will yield experiment-decision rule choices whose performance is quite satisfactory. From a practical point of view, this fact in itself is sufficient to justify its use.

IV. CONCLUSIONS

Well now, where do we stand with respect to justification of our approach to decision making under uncertainty? We have looked at alternate approaches in terms of axiomatic considerations, reasonableness arguments and comparison of average losses incurred. The conclusion that we now reach is that the objective Bayes (and its generalization the quasi-Bayes) approach is a good practical approach to this problem.

There are inconsistencies in this approach; however, the majority of them only occur when the inferred prior varies with the experiment. If the inferred prior is the same for all experiments under consideration, the justification arguments presented in Sec. I for the case of a single experiment apply to the quasi-Bayes approach.

REFERENCES

1. H. Robbins, "The Empirical Bayes Approach to Statistical Decision Problems," Ann. Math Stat., Vol. 35, March 1964, pp. 1-20.
2. D. V. Lindley, "The Use of Prior Probability Distributions in Statistical Inference and Decisions," in Proc. 4th Berkeley Symp. Mathematical Statistics and Probability. Berkeley, Calif.: Univ. of California Press, 1961, pp. 453-468.
3. A. Birnbaum, "Concepts of Statistical Evidence," in Philosophy, Science and Method, ed. S. Morgenbesser, P. Suppes and M. White, 1969.
4. R. D. Luce and H. Raiffa, Games and Decisions. New York: Wiley, 1957.
5. B. Shubert, "Bayesian Model of Decision Making as A Result of Learning From Experience," Ann. Math. Stat., Vol. 40, 1969, pp. 2127-2142.
6. B. DeFinetti, "Foresight, Its Logical Laws, Its Subjective Sources," in Studies in Subjective Probability, H. E. Kyburg, Jr. and H. Smokler, eds. New York: Wiley, 1963, pp. 97-158.
7. L. J. Savage, The Foundation of Statistics. New York: Wiley, 1954.
8. J. S. Maritz, Empirical Bayes Methods. New York: Barnes and Noble, 1970.
9. R. Radner and J. Marschak, "Note on Some Proposed Decision Criteria," in Decision Processes, R. Thrall, C. Coombs and R. Davis, eds. New York: Wiley, 1954, pp. 61-68.
10. T. L. Oberlin, Bayes Decision Rules Based on Objective Priors, Chapter 7, PhD. Thesis, Purdue University, Jan. 1972.